

CLAIMS

What is claimed is:

- 1 1. A host bus adapter for interconnecting a computer
2 system to a storage area network comprising
3 hardware for transmission of frames;
4 hardware for reception of frames;
5 memory for storage of frames;
6 a processor for processing frames, the processor
7 coupled to the hardware for transmission of frames,
8 the hardware for reception of frames, and the memory
9 for storage of frames;
10 wherein the processor is capable of inspecting
11 frames for encapsulated write requests and, if
12 encapsulated write request frames are found, de-
13 encapsulating the write request and forwarding the
14 write request through the hardware for transmission of
15 frames to a target node of the write request.
- 1 2. The host bus adapter of Claim 1, wherein the
2 processor further inspects frames for responses to
3 encapsulated write requests and, when such responses
4 are found, forwards the responses to a source node
5 from which the original encapsulated write request
6 frames came.
- 1 3. The host bus adapter of Claim 2, wherein the host
2 bus adapter is capable of simultaneous connection to
3 at least two links, having a transmitter and a
4 receiver for coupling to each link.
- 1 4. The host bus adapter of Claim 2, wherein the host
2 bus adapter is capable of responding to a query frame,
3 the query frame containing a request for status of any

4 path that might exist from the host bus adapter to a
5 specified target node.

1 5. The host bus adapter of Claim 1, wherein the host
2 bus adapter is capable of maintaining a mirrored
3 dataset on at least two target nodes.

1 6. The host bus adapter of Claim 5, wherein the host
2 bus adapter is capable of determining that it has no
3 direct path to a target node of the at least two
4 target nodes, and, when no direct path exists, is
5 capable of requesting that another node perform a
6 surrogate write to the target node.

1 7. The host bus adapter of Claim 5, wherein the host
2 bus adapter is capable of scanning nodes to determine
3 a node capable of performing a surrogate write to the
4 target node.

1 8. A node for connection to a storage area network
2 comprising
3 hardware for transmission of frames;
4 hardware for reception of frames;
5 memory;
6 at least one processor for processing frames, the
7 processor coupled to the hardware for transmission of
8 frames, the hardware for reception of frames, and the
9 memory for storage of frames;

10 wherein the processor is capable of inspecting
11 frames for encapsulated write requests and, if
12 encapsulated write request frames are found, de-
13 encapsulating the write request and forwarding the
14 write request through the hardware for transmission of
15 frames to a target node of the write request.

- 1 9. The node of Claim 8, wherein the processor
2 further inspects frames for responses to encapsulated
3 write requests and, when such responses are found,
4 forwards the responses to a source node from which the
5 original encapsulated write request frames came.
- 1 10. The node of Claim 9, wherein the node is capable
2 of simultaneous connection to at least two links,
3 having a transmitter and a receiver for coupling to
4 each link.
- 1 11. The node of Claim 9, wherein the node is capable
2 of responding to a query frame, the query frame
3 containing a request for status of any path existing
4 from the node to a specified target node.
- 1 12. The node of Claim 8, wherein the node is capable
2 of maintaining a mirrored dataset on at least two
3 target nodes.
- 1 13. The node of Claim 12, wherein the node is capable
2 of determining that it has no direct path to a target
3 node of the at least two target nodes, and, when no
4 direct path exists, is capable of requesting that
5 another node perform a surrogate write to the target
6 node.
- 1 14. The node of Claim 13, wherein the node is capable
2 of scanning nodes of a network to determine a node
3 capable of performing a surrogate write to the target
4 node.
- 1 15. A computer network comprising:
2 a first node;
3 a second node;
4 a first target node;

5 network interconnect providing communication
6 between the first node and the second node, and
7 providing communication between the second node and
8 the first target node;
9 wherein the second node is capable of inspecting
10 incoming frames for encapsulated write requests and,
11 if encapsulated write request frames are found, de-
12 encapsulating a write request from the encapsulated
13 write request frames and forwarding the write request
14 to the first target; and
15 wherein the second node further is capable of
16 inspecting frames received from the first target node
17 for responses to previously forwarded encapsulated
18 write requests and, when responses to previously
19 forwarded encapsulated write requests are found,
20 forwarding the responses to the first node.

1 16. The computer network of Claim 15, wherein the
2 second node is capable of responding to a path query
3 frame with a status of a path from the second node to
4 the first target node.

1 17. The computer network of Claim 16, wherein the
2 network interconnect is fibre channel compatible.

1 18. The computer network of Claim 16, further
2 comprising a second target node and the network
3 interconnect further provides communication between
4 the second target node and the first node; and
5 wherein the first node is capable of maintaining
6 a mirrored data set having a copy on the first target
7 node and the second target node.

1 19. A method of performing writes by a first node of
2 a storage area network to a mirrored dataset, the

3 dataset comprising a first copy on a first storage
4 node of the storage area network and a second copy on
5 a second node of the storage area network, the storage
6 area network having a surrogate-capable node, the
7 method comprising the steps of:
8 checking a status of a first path from the node
9 to the first storage node and of a second path from
10 the node to the second storage node;
11 the first path has good status and the second
12 path has bad status, then:
13 performing a write to the first copy of the
14 mirrored dataset over the first path;
15 polling the surrogate-capable node to determine
16 whether the surrogate-capable node has a path having
17 good status to the second storage node;
18 if the surrogate-capable node has a path having
19 good status to the second storage node, encapsulating
20 a write request to the second copy and transmitting
21 that encapsulated write request to the surrogate-
22 capable node; and
23 de-encapsulating the encapsulated write request
24 to the second copy and forwarding it from the
25 surrogate-capable node to the second storage node.

1 20. The method of Claim 19, further comprising the
2 steps of
3 transmitting a response to the write request to
4 the second copy from the second storage node to the
5 surrogate-capable node, and of
6 forwarding the response to the write request from
7 the surrogate-capable node to the first node.